

WHITE PAPER

Mainstreaming Server Virtualization: The Intel Approach

Sponsored by: Intel

John Humphreys

Tim Grieser

June 2006

IDC OPINION

The rapid spread of software-based server virtualization on Intel Architecture (IA) platforms over the past few years is making virtualization an important solution for scale-out hardware resource sharing. As virtualization is increasingly being used for production workloads, less complex and more efficient implementation strategies for deploying virtualization are being developed. With the uptake of the technology, it is increasingly important that integration occur between the hardware and virtualization layers in order to ensure the most efficient, high-performance, and reliable platform possible. The Intel approach is based on hardware-assisted virtualization, using the newly developed Intel® Virtualization Technology (Intel® VT), which provides specific hardware assists to enable virtual machine monitors (VMMs) to operate more efficiently.

IN THIS WHITE PAPER

In this IDC White Paper, we examine the growth of virtualization on Xeon® platforms and analyze the steps being taken by Intel to move virtualization into the mainstream on IA platforms.

INTRODUCTION

Virtualization in the form of partitioning a physical server into multiple virtual servers or virtual machines (VMs) is a long-established and widely accepted solution for hardware resource sharing on large, scale-up server platforms such as zOS and zVM mainframes as well as high-end Unix processors. In such environments, virtualization has evolved over many years to become a stable and robust technology for improving hardware utilization and efficiently processing production workloads, including high-value, mission-critical applications. The history of virtualization on such platforms includes evolution of the basic resource partitioning strategy from software-only implementation, to hardware-assisted partitioning software, to extended hardware-based virtualization support.


RAPID GROWTH OF VIRTUALIZATION

Virtualization has emerged during the past few years as an increasingly important resource-sharing strategy on scale-out Intel Architecture hardware platforms, and it is forecast to grow rapidly in the future, in terms of the number of servers supporting virtual machine capabilities. Today, virtualization on Intel is based largely on software implementations, such as VMMs or software hypervisors, which can partition physical servers into one or more "logical servers" or "virtual machines."

Each virtual machine can contain an image of an operating system (OS), such as Windows or Linux, which in turn can support standard applications. According to IDC, as shown in Figure 1, worldwide shipments of server hardware platforms with software virtualization capabilities are currently forecast to grow rapidly from 172,000 units in 2004 to 1,209,000 units in 2009, a compound annual growth rate (CAGR) of 47.7%. The number of virtualized system images running in partitions is expected to grow from 778,000 in 2004 to 5,107,000 in 2009, a CAGR of 45.7%. In the mix of virtualized server platforms, Windows and Linux images are expected to increase at CAGRs of 51.4% and 55.3%, respectively, whereas Unix images are forecast to increase at a lower CAGR of 34.7%.

FIGURE 1

Server Virtualization Shipments, 2004–2009



Source: IDC, 2006

INTEL PLATFORMS AND VIRTUALIZATION

Virtualization on high-end servers has been around for over 30 years, since IBM introduced it into mainframes. It transitioned from mainframe to RISC servers based on proprietary combinations of hardware and software (firmware, OS, hypervisor), and in recent years, we have seen virtualization emerge on Xeon processor-based servers with the advent of software-based virtualization solutions from vendors such as VMware and Microsoft, as well as the Xen open source project. Software virtualization on Intel has become a practical reality due also to the greatly increased server capacity and processor power that are now available on the Xeon processor-based platform.

Software Virtualization Approach

The implementation of partitioning and virtual machines on Intel Architecture platforms has been delivered through software VMMs, or hypervisors, most notably those from VMware (ESX Server), Microsoft (Microsoft Virtual Server), and Xen. In these software-only implementations, the software monitor provides a layer of abstraction (the virtual machine or logical server) between the base hardware (the Intel physical server) and the operating system (Windows, Linux) that allows each operating system running in a virtual machine to operate as if each OS has complete control of the underlying hardware. The VMM allows multiple virtual machines to exist concurrently on the same hardware platform, providing the ability to run multiple instances of operating systems and applications on a single physical server. Operating systems and applications running in virtual partitions are typically called "guests" to distinguish them from versions running in nonvirtualized server environments.

In essence, virtualization is a hardware resource-sharing strategy. The core advantage of implementing virtual machines is the ability to drive higher utilizations on physical servers by sharing the use of hardware resources across several workloads, moving away from the "one application per server" approach that often characterized the use of nonvirtualized servers in the past. Another advantage is the ability to isolate one workload from another, in terms of memory, data, and hard drive contents, avoiding such problems as driver conflicts when two applications are running on the same server.

Software-based virtual machine configurations are being used to assist software development and test environments. Common usage models include operating system and application migrations (running old and new versions of Windows on the same physical server), supporting multiple releases of an application on the same server, and providing multiple "logical" environments with different "virtual hardware" configurations for testing. Intel-based virtual machine configurations are increasingly being used to run production workloads such as in the case of server consolidations — running several production workloads on a single consolidated server, thereby achieving much higher average utilization while maintaining adequate headroom.

Limitations of Software Virtualization

While software virtualization offers many benefits, there are challenges in terms of overhead, performance, and supportability. These challenges arise because operating systems were designed to have full control over the underlying hardware and are not written to support resource sharing. New hardware virtualization capabilities, provided by Intel VT, in conjunction with today's software virtualization solutions, are designed to overcome these shortcomings.

THE INTEL VT APPROACH

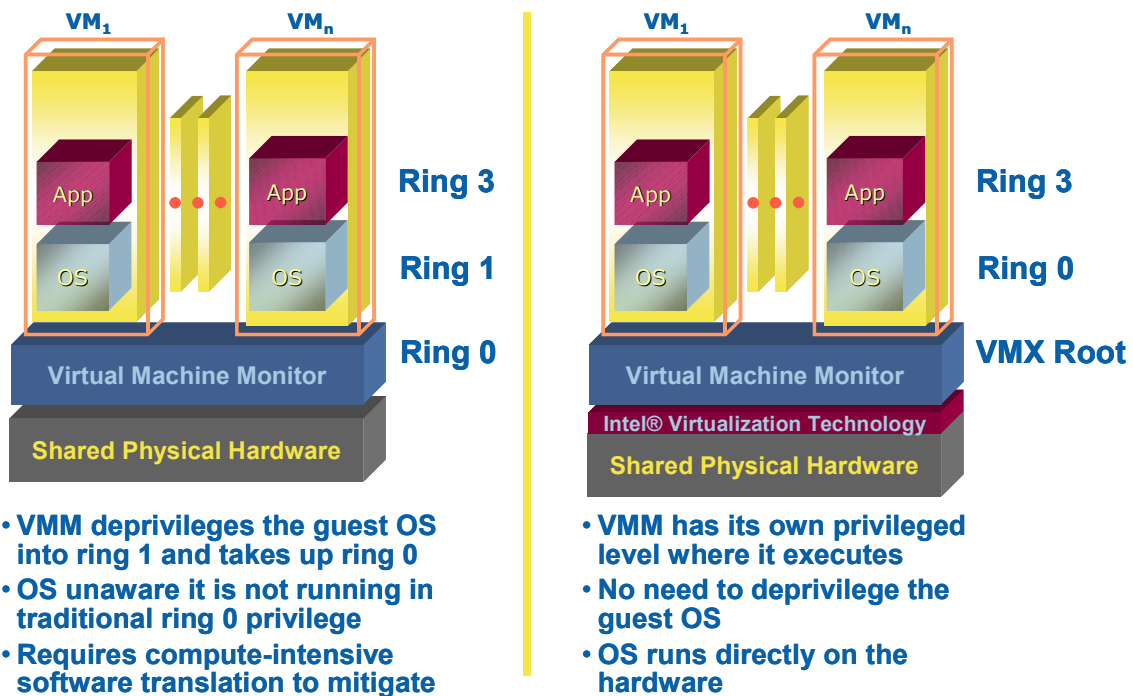
Intel has introduced hardware virtualization support through the new Intel Virtualization Technology now shipping on Intel processors. Intel VT is a set of hardware enhancements to Intel platforms that can improve the efficiency and capabilities of software virtualization solutions. The Intel Architecture is based on a ring privilege concept with four levels of privilege: ring 0 through ring 3. Ring 0 defines the highest privilege level and is dedicated to the operating system in a native environment.

The way virtualization is handled in a software-only mode is to run the VMM in ring 0, the most privileged layer, thereby depriving the operating systems to ring 1. While this approach works, it does present some challenges in terms of ring aliasing, nontrapping instructions, excessive faulting, CPU state context switching, and address space compression, which can lead to performance or reliability issues with the solution. Today, virtualization of current x86 CPUs requires complex software workarounds.

Going forward, VT is designed to eliminate these virtualization holes and the need for workarounds by running the VMM in a root layer, a new ring layer, thereby keeping the traditional ring structure of ring 0 for the operating system and ring 3 for the application software (see Figure 2).

FIGURE 2

Virtualization Technology: Pre- and Post-VT



Source: Intel, 2006

Intel VT provides the foundation for widely deploying virtualization solutions across a broad set of customer applications and production workload environments by working to address the challenges associated with software-only virtualization:

- ☒ **Overhead and performance.** Software-only VMMs introduce processing overheads to "emulate" server resources. These overheads occur in such areas as IO operations, memory management (paging), and simulating "privileged" CPU instructions. Intel VT, by supporting these functions with hardware virtualization, helps to reduce virtual machine overheads and hence expands the portfolio of applications and workloads suitable for virtualization.
- ☒ **Less complexity.** Software virtualization solutions today require the OS to go through the VMM when communicating with the underlying hardware because the OS is running in the address space and privilege level where applications are normally run. With software-only solutions, enabling an OS to run as a guest in a virtualized environment is achieved via binary translation and patching of the OS itself, to "trick" the OS into believing it is running on a "bare machine." Translation can be done dynamically, at runtime, or statically, in advance (known as paravirtualization). With the translation approach, every time the OS is updated (e.g., new release, service pack, patch), the VMM must also be patched to maintain support.

Hardware-assisted virtualization eliminates the need for binary translation or patching by providing a new architecture that allows the OS to run as an unmodified guest.

- ☒ **Reliability.** Many users worry that software translation or patching reduces the reliability of the overall solution. Intel VT addresses these concerns via hardware support for privilege ring expansion with resulting simplification of the hypervisor. Privilege ring expansion means that the VMM runs in a new, higher privilege ring, thus allowing the guest OS to run in its native privilege ring (ring 0).
- ☒ **Breadth of support.** Compared with today's software-only solutions, Intel VT broadens the number of OSs that can be supported. Intel VT enables direct support of unmodified, legacy operating systems. The aggressive ramp of Intel 64-bit Xeon solutions means that we will see an increased need for VMMs to support 64-bit guest OSs. Today's solutions do not support this capability, but solutions with Intel VT will. In addition, VT-based solutions support a full range of legacy OSs (multiple versions of Linux and Windows).
- ☒ **Flexibility.** A key goal of Intel VT is to make VMM software independent of specific OS software in order to remove the necessity of constantly updating VMMs to keep up with OS changes and patches.
- ☒ **Need for platform reliability.** As customers go from largely one-server/one-application environments today to many applications being hosted on a virtualized server, there is tremendous need for reliable hardware. In order to mitigate the risks associated with having many "eggs in one basket," users must search for the platform with the most comprehensive reliability, availability, and serviceability (RAS) features.

Overall, Intel VT hardware-assisted virtualization, combined with ongoing support from VMM software vendors, will provide enhanced virtual machine environments. The introduction of Intel VT marks the first step on Intel's long-term virtualization road map.

INDUSTRY SUPPORT FOR INTEL VT

The success of Intel VT depends on support from a broad ecosystem of system vendors, software vendors, channel partners, end users, and standards organizations. Intel is working with leading VMM software providers to deliver VT-enabled virtualization capabilities. Providers include Microsoft, VMware, and the open source community delivering open source VMM (Xen).

Intel is also actively working with leading platform vendors and channel partners to ensure that platforms ship with a BIOS that is compatible with Intel's latest processors that include Intel VT.

The goal is to provide the industry with a single-standard hardware-assisted virtualization platform that supports both paravirtualized and legacy guest operating systems, including both Xeon- and Itanium-based systems later this year as well as client devices.

Intel has announced support for defined VMM interfaces to the host operating system, the system management software stack, and/or platform firmware and the platform devices (see Figure 3). This interface helps to offload low-level activities to minimize the overhead associated with running the virtualization layer and to increase the robustness, security, and reliability of the virtualized services. Intel VT also allows VMM on VT platforms to be executed directly on a fully Intel-verified IA-32 architecture and instruction set.

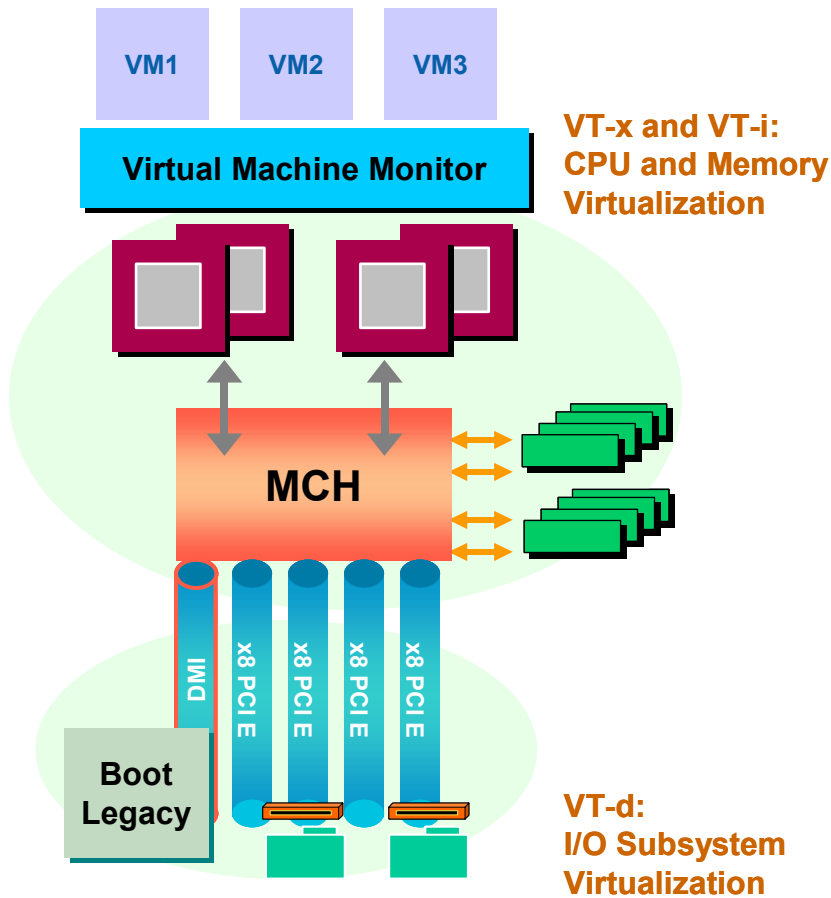
The release of CPU virtualization in late 2005 was the first step in a series of Intel virtualization solutions under the umbrella name of VT. Intel has focused its current development efforts on assisting virtualization of processors and memory. Going forward, the company is taking the concepts to the I/O subsystems to further improve the virtualization capabilities of the entire platform. In March 2006, Intel released a preliminary specification for Intel Virtualization Technology for Directed I/O, or VT-d, which it had previously disclosed with relevant platform vendors.

VT-d provides hardware assists in the chip set for direct memory access (DMA) remapping and direct assignment. By supporting DMA remapping in hardware, VT-d provides increased memory protection from errant DMA code and improved system reliability. VMMs can also use VT-d to direct-assign I/O resources to specific virtual machines. Direct assignment provides the unmodified guest OS with direct access to modern physical device functions for which emulated drivers may not be written by the VMM vendor. It also improves I/O performance for specific usage models like mission-critical or test and development environments. Intel is enabling VMM vendors with VT-d silicon in 2006 and will begin to ship products in line with ecosystem readiness.

Further out, Intel is also investigating more granular memory virtualization to further improve VMM performance and robustness.

FIGURE 3

Intel's Support for Hardware Virtualization



Source: Intel, 2006

FUTURE OUTLOOK

The outlook for virtualization is incredibly strong. Recent survey work by IDC suggests that one-fifth of the server installed base of companies that already employ virtualization technologies is virtualized. The expectation is that these organizations will increase that percentage to nearly 50% by the end of 2006.

There is growing evidence that the technology is moving beyond early adopters and is gaining acceptance among more mainstream customers as the market finds additional uses for virtualization and a wide range of benefits, as detailed in the following sections.

Server Consolidation

The ability to run multiple partitions with each partition includes virtual BIOS, OS, and applications. Software-only virtualization provides sharing of physical hardware resources with complete software isolation. It leads to reduced hardware costs (though software costs will be similar), easier administration and simplified manageability, as well as reduced power consumption, cabling, and costly datacenter real estate. This is the most popular current usage for virtualization.

- ☒ Consolidation of IT infrastructure
- ☒ Datacenter in a box
- ☒ Mixed OS environment

Test and Development

An increasingly popular use is to run pilot solutions in different virtual partitions. This approach realizes all the benefits of server consolidation (reducing the number of "nonproductive" servers) and allows increased agility through accelerated application rollout. New applications can be validated inside partitions before being rolled out across the datacenter, increasing the robustness of the solution.

- ☒ Improved test hardware utilization
- ☒ New application deployment
- ☒ Rapid deployment

Hardware Migrations and Upgrades

The benefits that virtualization offers to hardware migrations are similar to the benefits it offers to test and development environments. Virtualization allows a staged migration of OS and applications onto new hardware. By validating the solutions in virtual partitions, IT shops reduce the possibility of disruption of service due to migration. When migrating to a new architecture, staging the migration in virtualized servers allows robust validation before a widespread rollout.

- ☒ Change management
- ☒ Operating system migrations
- ☒ Patches and upgrades

Business Continuity

Most IT shops today deploy some form of failover, usually involving replicated servers. Because many failures are associated with software, users are recognizing that replicating servers in virtual partitions on the same machine provides the benefits of failover at a reduced hardware cost. This replication, of course, necessitates that the server itself be extremely reliable, accounting for the popularity of Intel Xeon

processor MP-based servers with their advanced RAS features. As the latest Xeon and Itanium servers continue to offer enhanced reliability, this usage will increase.

- ☒ Failover of virtual machine
- ☒ IT partition
- ☒ Disaster recovery (backups)

Capacity Planning

Virtual partitions can be sized and resized as required. This capability offers great flexibility to IT shops, which can configure large compute resources when required (e.g., close of the quarter), then scale down as needed. By allocating compute resources as needed, IT shops can optimize overall server utilization.

In the future, provisioning of compute resources will become even more dynamic, providing the ability to add and remove resources as required.

- ☒ Infrastructure pooling
- ☒ Dynamic application scaling
- ☒ Utility computing

Load Balancing

Load balancing can be attained when policy-based tools monitor utilization and virtual machines are moved as needed to balance peak capacity with headroom.

- ☒ Workload balancing across platforms
- ☒ Headroom planning
- ☒ Dynamic application scaling

The use cases for virtualization have expanded dramatically during the time the technology has been in the market. Initially, virtualization was seen as a tool to improve hardware utilization through consolidation. Now it is also a means to extend application life cycles by rehosting older applications from legacy operating systems onto more powerful hardware.

The application of virtualization has now expanded into business continuity, capacity planning, and a foundation for policy-based automation — the embryonic stages of true utility computing.

This ever-expanding portfolio of usage for virtualization is what makes the technology so compelling to end users and, in turn, is why Intel, its hardware partners, and leading virtualization vendors are working together to create a robust platform that encompasses the hardware, software, and management layers of the solution. This collaborative approach is critical if the challenges of virtualization are to be overcome and the full potential of the technology is to be realized.

CHALLENGES/OPPORTUNITIES

The challenges for Intel around hardware-assisted virtualization center primarily on driving acceptance of the solution. Intel will need to continue to demonstrate to both customers and software virtualization partners how hardware-assisted virtualization can improve virtualization performance and expand the use cases and applications that can be run successfully on the platform.

Intel will need to continue to develop the ecosystem around Intel VT and foster an environment in which both software vendors and Intel can partner effectively. The first challenge the company is undertaking is orchestrating VMM support for Intel VT with the software suppliers: VMware, Microsoft, and Xen.

This rollout of the VT hardware and supporting monitors is being accomplished by focusing on the needs of end users. Success in this rollout would lead to increased credibility for the solution and would assist Intel in staking out a winning position in the virtualization of x86 systems.

In addition, Intel will continue to remain vigilant as it differentiates itself based on hardware virtualization capabilities, continues to innovate, and promotes standards. Innovating in the industry-standard market, where partners play a key role in any vendor's success, is a game that Intel has long learned to play. The key will be for the company to continue applying these lessons to the virtualization market, ensuring a place for software partners to enhance the capabilities of the products they provide.

CONCLUSION

Virtualization is shaping up to be one of the major trends that impacts the server market as well as customer datacenters. The technology is already finding a broad set of usages, ranging from improved hardware utilization in test and development infrastructure, to application life-cycle extension, to high availability and disaster recovery.

The emergence of virtualization as a horizontal enabling technology is a major reason customers are so bullish on the solution. However, for virtualization to reach its full potential and broaden the portfolio of applications it can run efficiently, tighter integration with hardware is necessary.

Intel is making the investments needed to allow more widespread and mainstream adoption by working to incorporate VT in the entire range of the company's platforms. This move helps to ensure the availability of efficient, robust, and high-performance platforms running a wide range of virtualization technologies.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2006 IDC. Reproduction without written permission is completely forbidden.